Chinese Restaurant Process for cognate clustering: A threshold free approach

Taraka Rama
Institute of Linguistics
University of Tübingen
Germany
taraka-rama.kasicheyanula@uni-tuebingen.de

arXiv:1610.06053v1 [cs.CL] 19 Oct 2016

Abstract

In this paper, we introduce a threshold free approach, motivated from Chinese Restaurant Process, for the purpose of cognate clustering. We show that our approach yields similar results to a linguistically motivated cognate clustering system known as LexStat. Our Chinese Restaurant Process system is fast and does not require any threshold and can be applied to any language family of the world.

1 Introduction

Identification of cognates is an important task while establishing genetic relations between languages that are hypothesized to have descended from a single language in the past. For instance, English *hound* and German *Hund* are cognates which can be traced back to Proto-Germanic stage. Highly accurate automatic identification of cognates is desired for reducing the effort required in analyzing large language families such as Indo-European (Bouckaert et al., 2012) and Austronesian (Greenhill and Gray, 2009) can take up decades of effort, when performed by hand. A automatic cognate identification system can be helpful for historical linguists to analyze supposedly related language families and also fasten up the making of cognate databases which can be then be analyzed using Bayesian phylogenetic methods (Atkinson and Gray, 2006).

In this paper, we work with Swadesh word lists of multiple language groups and attempt to cluster related words together using a non-parametric process known as Chinese Restaurant Process (Gershman and Blei, 2012). We use the sound similarity matrix trained in an unsupervised fashion (Jäger, 2013) for the purpose of computing similarity between two words. The CRP based algorithm is similar to the CRP variant of the K-means algorithm introduced by Kulis and Jordan (2011). Our CRP algorithm does not require any threshold and only has a single hyperparameter known as α which allows new clusters to be formed without the requirement of threshold or the number of clusters to be known beforehand.

Previous work by List et al. (2016) and Hauer and Kondrak (2011) employ a hand crafted or a machine learned word similarity measure to compute pair-wise distances between words. The pair-wise distance matrix is then supplied to a clustering algorithm such as average linkage clustering (Manning and Schütze, 1999) for inferring a tree structure of the words. The average linkage clustering algorithm is an agglomerative algorithm that merges individual clusters until a single cluster is left. The clustering process can be interrupted if the average similarity between two clusters falls below a predetermined threshold.

The agglomerative algorithm is simple and usually yields reasonable results across various language families (List, 2012). However, the method suffers from a major drawback that the threshold needs to be known beforehand for achieving high accuracy. In a recent paper, List et al. (2016) use a clustering algorithm known as InfoMap for the purpose of clustering cognates in Sino-Tibetan language groups. The InfoMap algorithm also requires a threshold for finding cognates. The authors find that the algorithm works well if the threshold is adjusted across language groups. In this paper, we compare our system against the LexStat system and show that our system yields comparable results.

The structure of the paper is as followed. We define the cognate clustering problem in section 2. In section 3, we describe the string alignment algorithm for the purpose of computing similarity between two strings. We describe our CRP algorithm in section 4. We describe the evaluation of our experiments and datasets in section 5. We present

the results of our experiments and discuss them in section 6. We conclude the paper in section 7.

2 Cognate clustering

The phylogenetic inference methods require cognate judgments which are only available for a small number of well-studied language families such as Indo-European and Austronesian. For instance, the ASJP database (Brown et al., 2013)[1] provides Swadesh word lists (of 40 length that are supposedly important for identifying genetic relationships between languages) transcribed in a uniform format for more than 60% of the world's languages.[2] An example of such a word list is given below:

	ALL	AND	ANIMAL	...
English	ol	End	Enim3l	...
German	al3	unt	tia	...
French	tu	e	animal	...
Spanish	to8o	i	animal	...
Swedish	ala	ok	y3r	...

Table 1: Example of a word list, in ASJP transcription for five languages belonging to Germanic (English, German, and Swedish) and Romance (Spanish and French) subfamilies.

The task at hand is to automatically cluster words according to genealogical relationship. This is achieved by computing similarities between all the word pairs belonging to a meaning and then supplying the resulting distance matrix as an input to a clustering algorithm. The clustering algorithm groups the words into clusters by optimizing a similarity criterion. The similarity between a word pair can be computed using supervised approaches (Hauer and Kondrak, 2011) or by using sequence alignment algorithms such as Needleman-Wunsch (Needleman and Wunsch, 1970) or Levenshtein distance (Levenshtein, 1966). An example of a pairwise distance matrix for meaning "all" is shown in table 2.

3 Sequence alignment

The Needleman-Wunsch algorithm is the similarity counterpart of the Levenshtein distance. The Needleman-Wunsch algorithm maximizes similarity whereas Levenshtein distance minimizes the distance. In the Needleman-Wunsch algorithm, a character or sound segment match increases the similarity by 1 and a character mismatch has a weight of -1. In contrast to Levenshtein distance which treats insertion, deletion, and substitution equally, the Needleman-Wunsch algorithm introduces a gap opening (deletion operation) penalty parameter that has to be learned separately. A second parameter known as gap extension penalty has lesser penalty than the gap opening parameter and models the fact that deletions occur in chunks (Jäger, 2013).

	ol	al3	tu	to8o	ala
ol		0.28	0.99	0.99	0.4
al3	0.28		0.94	0.99	0.01
tu	0.99	0.94		0.55	0.99
to8o	0.99	0.99	0.55		0.99
ala	0.4	0.01	0.99	0.99	

Table 2: An example of a pairwise distance matrix between all the words for meaning "all".

The (vanilla) Needleman-Wunsch algorithm is not sensitive to segment pairs and a realistic algorithm should assign high similarity between sound correspondences such as /s/ $\sim$ /h/ than the sound pair /p/ $\sim$ /r/.

In dialectology (Wieling and Nerbonne, 2015), similarity between two segments is estimated using PMI. The PMI score of two sounds i and j is defined as followed:

$$\text{PMI}(i,j) = \log(\frac{p(i,j)}{q(i) \cdot q(j)}) \qquad (1)$$

where, $p(i,j)$ is the relative frequency of i,j occurring at the same position in th aligned word pairs whereas, $q(.)$ is the relative frequency of a sound in the whole word list. A positive PMI value indicates that a segment pair cooccurs together whereas, a negative PMI value indicates lack of cooccurrence. This can be interpreted as a strength of relatedness between two segments.

In this paper, we use the PMI matrix (of ASJP sound segments) inferred by Jäger (2013) for computing the similarity between a word pair. Jäger (2013) shows that the PMI matrix shows positive weights for sound pairs such as /p/ $\sim$ /b/, /t/ $\sim$ /d/, and /s/ $\sim$ /h/.

4 CRP

In this section, we describe the CRP algorithm and motivate its suitability for cognate clustering.

[1] asjp.clld.org

[2] However, the cognacy judgments are only available for a subset of language families.

Given a meaning M and the word similarity matrix S of dimensions $N \times N$, the CRP algorithm works as follows. The CRP outputs K clusters and the clustering $l_1, \ldots l_K$.

1. Initially, assign a word w_n to l_n where $K = N$.
2. Repeat until convergence:
   - For each word w_n:
     - Remove w_n from its cluster.
     - Compute s_{nk} the average similarity of w_n to all words in cluster k.
     - If $\arg\max_k s_{nk} < \alpha$ assign w_n to a new cluster.
     - Else, assign w_n to the cluster k where $k = \arg\max_k s_{nk}$.

The current algorithm uses the criterion of average similarity to assign a word to a cluster. A word is assigned to the cluster with which it exhibits the highest average similarity. The intuition behind this decision is that the word should, on an average, be similar to the rest of the words in a cluster. [3]

The magnitude of the α parameter determines the number of new clusters. A value of 0.01 is sufficient for the purpose of forming new clusters. The word similarity is always non-negative and we use a ReLU transformation ($max(0, x)$) that transforms negative similarity scores to 0. The CRP algorithm identifies cognate clusters of uneven sizes and can also form singleton clusters due to the simple initialization. In our experiments, we find that three full scans of the data are sufficient for the algorithm to reach a local maximum.

5 Experiments

Baseline We use a vanilla Needleman-Wunsch with a gap opening penalty of -1 and a gap extension penalty of -0.5 as the baseline in our experiments.

LexStat LexStat (List, 2012) is a system offering state-of-the-art alignment algorithms for aligning word pairs and clustering them into cognate sets. The LexStat system weighs matches between sounds using a handcrafted segment similarity matrix that is informed by historical linguistic literature.

[3]The average similarity criterion can be modified to the maximum similarity criterion. This is commonly known as single linkage clustering (Manning and Schütze, 1999).

5.1 Evaluation

We evaluate the results of clustering analysis using B-cubed F-score (Amigó et al., 2009). The B-cubed scores are defined for each individual item as followed. The precision for an item is defined as the ratio between the number of cognates in its cluster to the total number of items in its cluster. The recall for an item is defined as the ratio between the number of cognates in its cluster to the total number of expert labeled cognates. The B-cubed precision and recall are defined as the average of the items' precision and recall across all the clusters. Finally, the B-cubed F-score for a meaning, is computed as the harmonic mean of the average items' precision and recall. The B-cubed F-score for the whole dataset is given as the average of the B-cubed F-scores across all the meanings.

Both Hauer and Kondrak (2011) and List et al. (2016) use B-cubed F-scores to test their cognate clustering systems.

5.2 Datasets

IELex database The Indo-European Lexical database was created by Dyen et al. (1992) and curated by Michael Dunn. The IELex database is not transcribed in uniform IPA and retains many forms transcribed in the Romanized IPA format of Dyen et al. (1992). We cleaned the IELex database of any non-IPA-like transcriptions and converted the cleaned subset of the database into ASJP format. The cleaned subset has 52 languages and 210 meanings.

Austronesian vocabulary database The Austronesian Vocabulary Database (ABVD) (Greenhill and Gray, 2009) has word lists for 210 Swadesh concepts and 378 languages.[4] The database does not have transcriptions in a uniform IPA format. We removed all symbols that do not appear in the standard IPA and converted the lexical items to ASJP format. For comparison purpose, we use randomly selected 100 languages' dataset in this paper.[5]

Short word lists with cognacy judgments Wichmann and Holman (2013) and List (2014) compiled cognacy wordlists for subsets of families from various scholarly sources such as comparative handbooks and historical linguistics' articles. The details of this compilation is given below. For

[4]`http://language.psy.auckland.ac.nz/austronesian/`

[5]LexStat takes many hours to run on a dataset of 100 languages.

each dataset, we give the number of languages/the number of meanings in parantheses.

- Wichmann and Holman (2013): Afrasian (21/40), Mayan (30/100), Mixe-Zoque (10/100), Mon-Khmer (16/100).
- List (2014): ObUgrian (21/110; Hungarian excluded from Ugric sub-family).

6 Results

The B-cubed F-scores of different systems are shown in table 3. The CRP based PMI system performs better than LexStat on four datasets. The CRP algorithm performs slightly worse than the LexStat system on Austronesian and Indo-European language families by two points. The LexStat performs better than the PMI-CRP system only on the Ugric languages dataset. The LexStat system's clustering threshold has been tuned on many smaller datasets whereas, the PMI-CRP does not require any tuning of the threshold and comes closer or performs better than the LexStat system. We also provide the average of B-cubed F-scores across different datasets. The results show that the PMI-CRP system is close to the performance of the LexStat system.

Dataset	LexStat	NW-CRP	PMI-CRP
Afrasian	76.54	78.2	**81.22**
Austronesian	**74.9**	71.89	72.39
Mayan	78.6	80.38	**80.75**
Mixe-Zoque	91.45	88.61	**92.35**
Indo-European	**77.56**	67.2	74.89
Mon-Khmer	80.49	78.11	**81.69**
ObUgrian	**92.19**	73.4	86.78
Average	81.68	76.83	81.44

Table 3: Average B-cubed F-scores of different systems. The suffix "-CRP" stands for the CRP algorithm applied to Needleman-Wunsch (NW) and PMI word similarity methods.

6.1 Match between predicted and obtained clusters

We examine the match between the number of predicted clusters and the number of true clusters for the PMI-CRP system across meanings. We report the correlations in table 4. The correlations suggest that the number of predicted clusters correlate highly with the true number of clusters across datasets.

Dataset	PMI-CRP
Afrasian	72.69
Austronesian	70.66
Mayan	81.35
Mixe-Zoque	88.94
Indo-European	76.41
Mon-Khmer	82.11
Ugric	73.82

Table 4: Pearson's correlation between the true number of clusters and the number of predicted clusters across language families.

6.2 Error analysis

In the case of Indo-European, the PMI-CRP system fails to group all the reflexes for the meaning "five", "fingernail", "three", "two", and "name" into a single cognate cluster. The reason for this behaviour is the extensive phonological change that affected cognates across the daughter subgroups. The LexStat system also shows similar behaviour when the true number of cognate clusters is 1.

7 Conclusion

In this paper, we introduced a CRP based clustering algorithm that is threshold free. The program takes less than two minutes for clustering a large dataset of 100 languages such as Austronesian. We tested the algorithm on a wide range of language families and showed the algorithm yields close or better results than LexStat. Based on the results, we claim that the algorithm can be useful for the comparative linguists to analyze putative language relations at a quick pace.

The main limitation of the algorithm is that it fails to retrieve clusters for meanings such as "what", "who", and "we" (in Indo-European) which show high phonological divergence. In comparison, even LexStat makes mistakes when clustering these meanings. Whenever the reflexes show similar word forms, in the case of Mayan (meanings: "water" and "die"), the algorithm groups all the reflexes into a single cluster without any error.

As part of future work, we plan to use the CRP algorithm for clustering meanings across different language families available in the ASJP database and then supply the cognate clusters to a Bayesian phylogenetic inference software such as MrBayes (Ronquist and Huelsenbeck, 2003) for inferring Bayesian trees for the languages of the world.

References

[Amigó et al.2009] Enrique Amigó, Julio Gonzalo, Javier Artiles, and Felisa Verdejo. 2009. A comparison of extrinsic clustering evaluation metrics based on formal constraints. *Information retrieval*, 12(4):461–486.

[Atkinson and Gray2006] Quentin D. Atkinson and Russell D. Gray. 2006. How old is the Indo-European language family? Progress or more moths to the flame. In Peter Forster and Collin Renfrew, editors, *Phylogenetic Methods and the Prehistory of Languages*, pages 91–109. The McDonald Institute for Archaelogical Research, Cambridge.

[Bouckaert et al.2012] Remco Bouckaert, Philippe Lemey, Michael Dunn, Simon J. Greenhill, Alexander V. Alekseyenko, Alexei J. Drummond, Russell D. Gray, Marc A. Suchard, and Quentin D. Atkinson. 2012. Mapping the origins and expansion of the Indo-European language family. *Science*, 337(6097):957–960.

[Brown et al.2013] Cecil H. Brown, Eric W. Holman, and Søren Wichmann. 2013. Sound correspondences in the world's languages. *Language*, 89(1):4–29.

[Dyen et al.1992] Isidore Dyen, Joseph B. Kruskal, and Paul Black. 1992. An Indo-European classification: A lexicostatistical experiment. *Transactions of the American Philosophical Society*, 82(5):1–132.

[Gershman and Blei2012] Samuel J Gershman and David M Blei. 2012. A tutorial on bayesian nonparametric models. *Journal of Mathematical Psychology*, 56(1):1–12.

[Greenhill and Gray2009] Simon J. Greenhill and Russell D. Gray. 2009. Austronesian language phylogenies: Myths and misconceptions about Bayesian computational methods. *Austronesian Historical Linguistics and Culture History: A Festschrift for Robert Blust*, pages 375–397.

[Hauer and Kondrak2011] Bradley Hauer and Grzegorz Kondrak. 2011. Clustering semantically equivalent words into cognate sets in multilingual lists. In *Proceedings of 5th International Joint Conference on Natural Language Processing*, pages 865–873, Chiang Mai, Thailand, November. Asian Federation of Natural Language Processing.

[Jäger2013] Gerhard Jäger. 2013. Phylogenetic inference from word lists using weighted alignment with empirically determined weights. *Language Dynamics and Change*, 3(2):245–291.

[Kulis and Jordan2011] Brian Kulis and Michael I Jordan. 2011. Revisiting k-means: New algorithms via bayesian nonparametrics. *arXiv preprint arXiv:1111.0352*.

[Levenshtein1966] Vladimir I. Levenshtein. 1966. Binary codes capable of correcting deletions, insertions and reversals. In *Soviet physics doklady*, volume 10, page 707.

[List et al.2016] Johann-Mattis List, Philippe Lopez, and Eric Bapteste. 2016. Using sequence similarity networks to identify partial cognates in multilingual wordlists. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 599–605, Berlin, Germany, August. Association for Computational Linguistics.

[List2012] Johann-Mattis List. 2012. LexStat: Automatic detection of cognates in multilingual wordlists. In *Proceedings of the EACL 2012 Joint Workshop of LINGVIS & UNCLH*, pages 117–125, Avignon, France, April. Association for Computational Linguistics.

[List2014] J.-M. List. 2014. *Sequence comparison in historical linguistics*. Düsseldorf University Press, Düsseldorf.

[Manning and Schütze1999] Christopher D Manning and Hinrich Schütze. 1999. *Foundations of statistical Natural Language Processing*. MIT Press.

[Needleman and Wunsch1970] Saul B. Needleman and Christian D. Wunsch. 1970. A general method applicable to the search for similarities in the amino acid sequence of two proteins. *Journal of Molecular Biology*, 48(3):443–453.

[Ronquist and Huelsenbeck2003] Fredrik Ronquist and John P Huelsenbeck. 2003. Mrbayes 3: Bayesian phylogenetic inference under mixed models. *Bioinformatics*, 19(12):1572–1574.

[Wichmann and Holman2013] Søren Wichmann and Eric W Holman. 2013. Languages with longer words have more lexical change. In *Approaches to Measuring Linguistic Differences*, pages 249–281. Mouton de Gruyter.

[Wieling and Nerbonne2015] Martijn Wieling and John Nerbonne. 2015. Advances in dialectometry. *Annu. Rev. Linguist.*, 1(1):243–264.